WHEN SHAPE MATTERS: DEFORMATIONS OF TILING SPACES

ALEX CLARK AND LORENZO SADUN

Abstract. We investigate the dynamics of tiling dynamical systems and their deformations. If two tiling systems have identical combinatorics, then the tiling spaces are homeomorphic, but their dynamical properties may differ. There is a natural map $\mathcal{I}$ from the parameter space of possible shapes of tiles to H^1 of a model tiling space, with values in $\mathbb{R}^d$. Two tiling spaces that have the same image under $\mathcal{I}$ are mutually locally derivable (MLD). When the difference of the images is "asymptotically negligible", then the tiling dynamics are topologically conjugate, but generally not MLD. For substitution tilings, we give a simple test for a cohomology class to be asymptotically negligible, and show that infinitesimal deformations of shape result in topologically conjugate dynamics only when the change in the image of $\mathcal{I}$ is asymptotically negligible. Finally, we give criteria for a (deformed) substitution tiling space to be topologically weakly mixing.

1991 *Mathematics Subject Classification.* 37B50, 52C23, 37A20, 37A25, 52C22.

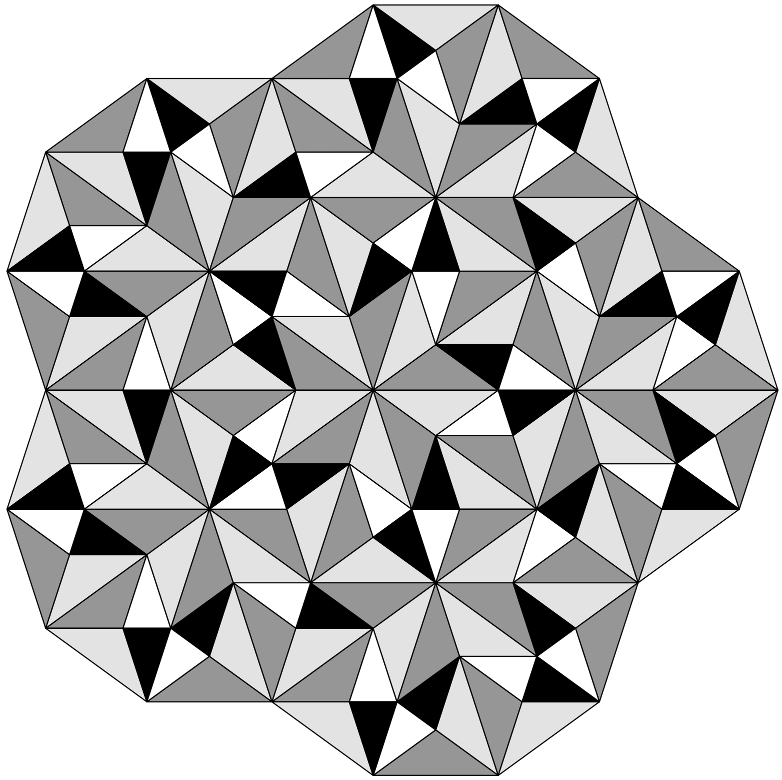

FIGURE 1. A patch of a Penrose tiling

1. Introduction and Statement of Results

A tiling is described by a combination of combinatorial data (which tiles meet which others) and by geometric data (the shape and location of each tile). Tilings with the same combinatorics may have different geometry. Compare figure 1, which shows a patch of a Penrose tiling, to figure 2, which shows a corresponding patch of a combinatorially identical but geometrically different tiling. The translation group acts on the continuous hulls of both tilings. How do the dynamics compare?

To avoid trivialities, we shall only consider *nonperiodic* tilings in this paper. Moreover, we shall assume that the tiles are polyhedra that meet full-face to full-face, and that there are only a finite set of tile types (a.k.a. prototiles) up to translation. Although these conditions may seem to be restrictive, any tiling that has finite local complexity with respect to translations is mutually locally derivable (MLD, see definition below) to a tiling meeting these conditions, via derived Voronoi tilings [Pr]. We also assume that our tiling spaces are minimal. This is equivalent to the condition that each patch of each tiling appear in all other tilings with bounded gaps.

Deformations of tiling spaces were considered in [SW]. As detailed there, to describe the shape and size of each tile, one must specify the displacement vector that corresponds to each edge of each prototile. If, somewhere in the tiling, translations of two prototiles

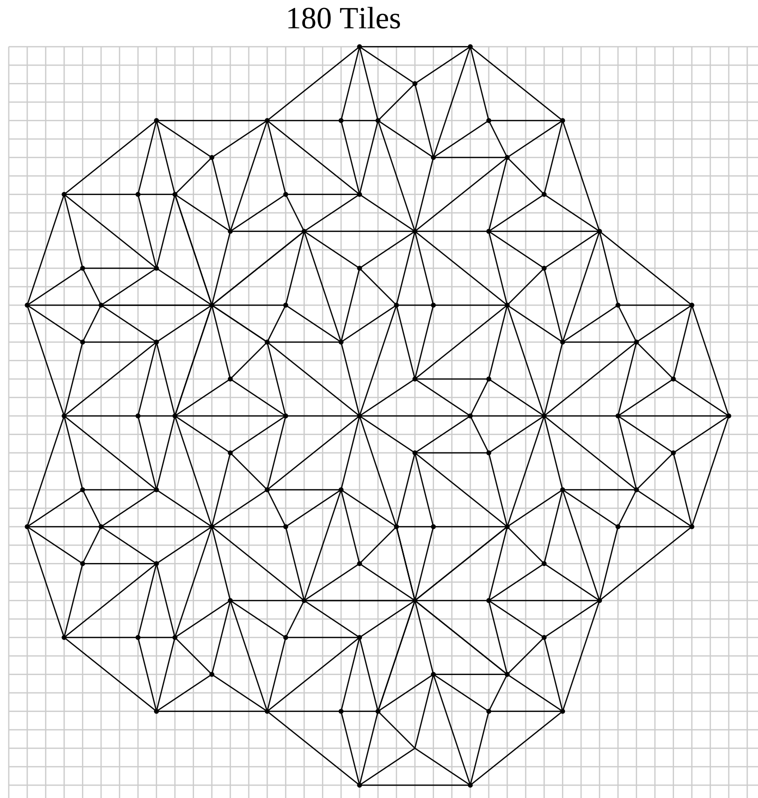

FIGURE 2. A corresponding patch of a deformed Penrose tiling

meet along a common edge, then those two edges must have the same associated displacement vector. Moreover, the sum of the displacement vectors around each tile must be zero. Each deformation thus corresponds to a solution of a homogeneous linear system of equations subject to additional conditions which are open in the solution space: the sequence of edges around a tile does not cross itself and has winding number $+1$ and bounds a 2-cell of a tile. If $d > 2$ we also require the edges of a 2-cell to be coplanar, and in general for all the edges of a k-cell of a tile to lie in a k-dimensional hyperplane of $\mathbb{R}^d$.

Put another way, the shapes of the prototiles are determined by a function f from the set of prototile edges (modulo certain identifications) to $\mathbb{R}^d$. This function may be viewed as a 1-cochain on the CW-complex obtained by taking the disjoint union of all the prototiles, modulo the identification of edges (and other structures of dimension less than d) where tiles can meet. This is precisely the complex Γ considered by Anderson and Putnam in their computation of topological invariants of substitution tiling spaces [AP]. (The Anderson-Putnam complex is defined for all translationally finite tilings, although it has been applied primarily to substitution tilings. See [G] for applications in a more general context.) That is, $f \in C^1(\Gamma, \mathbb{R}^d)$, and we may think of "shape space" as a subset of $C^1(\Gamma, \mathbb{R}^d)$.

We denote vectors by bold-face roman letters, tiling spaces by $\mathcal{T}$ with appropriate subscripts, individual tilings by roman letters from the end of the alphabet (usually x or

y), and shape parameters by f or g. The translate of the tiling x by $\mathbf{z} \in \mathbb{R}^d$ is denoted $x + \mathbf{z}$.

Fix a tiling space $\mathcal{T}$, and consider deformations of this space. For any choice f of edge vectors meeting the above requirements, we can consider tilings whose tiles are described by these edge vectors, but whose combinatorics (which tiles meet and how) are the same as the original tilings in $\mathcal{T}$. Let $\mathcal{T}_f$ denote the space of such deformed tilings, on which $\mathbb{R}^d$ acts naturally by translation. The primary focus of this paper is the extent to which the dynamical system $\mathcal{T}_f$ depends on the function f. If all we care about is the topological space, a theorem of [SW] shows that it doesn't:

Theorem 1.1. [SW] *Let $f, g \in C^1(\Gamma, \mathbb{R}^d)$ be admissible shape functions for a fixed tiling space. Then $\mathcal{T}_f$ and $\mathcal{T}_g$ are homeomorphic.*

But what about dynamics? In a previous paper [CS] we showed how 1-dimensional substitution tiling space dynamics depend on the lengths of the tiles. Here we extend that analysis to higher dimensions, and also recast our 1-dimensional results in terms of topological invariants. One natural question is when two tiling spaces have topologically conjugate dynamics. Another is when they have dynamics that are intertwined by a local map. This involves the notion of *mutual local derivability*, first introduced in [BSJ].

Let x and y be two tilings, possibly with different set of tiles. The tiling y is said to be *locally derivable* from x if there exists a length R such that, if $\mathbf{z}_1, \mathbf{z}_2 \in \mathbb{R}^d$ and $x - \mathbf{z}_1$ agrees with $x - \mathbf{z}_2$ on a ball of radius R around the origin, then $y - \mathbf{z}_1$ agrees with $y - \mathbf{z}_2$ on a ball of radius 1 around the origin. In other words, the type and exact placement of the tile at a point $\mathbf{z}$ in y depends only on the patch of radius R around $\mathbf{z}$ in x. If y is locally derivable from x and x is locally derivable from y, then x and y are said to be mutually locally derivable (MLD) tilings.

The MLD concept extends to tiling spaces. If x and y are MLD tilings, then the closure of the translational orbit of x is topologically conjugate to the closure of the translation orbit of y, via a conjugacy that takes x to y, and thus depends only on local data. More generally, we say that two tiling spaces $\mathcal{T}$ and $\mathcal{T}'$ are MLD if there exists a topological conjugacy $\phi : \mathcal{T} \mapsto \mathcal{T}'$ such that, for some R, the patch of size 1 around the origin in $\phi(x)$ can be determined exactly from the patch of size R around the origin in x, and the patch of size 1 in x can be determined exactly from the patch of size R in $\phi(x)$. This is a natural generalization of the concept of "sliding block codes" for subshifts. However, while all continuous maps of subshifts are sliding block codes (see, e.g., [LM]), there exist topologically conjugate tiling spaces that are not MLD [RS, Pe, CS].

In Section 2, we construct a natural map $\mathcal{I}$ from the space of shape parameters to the Cech cohomology group $H^1(\mathcal{T}, \mathbb{R}^d)$ and show that the dynamical properties of $\mathcal{T}_f$ depend only on $\mathcal{I}(f)$:

Theorem 1.2 (Theorem 2.1)**.** *Suppose f and g are admissible shape parameters with $\mathcal{I}(f) = \mathcal{I}(g)$. Then $\mathcal{T}_f$ and $\mathcal{T}_g$ are mutually locally derivable.*

Furthermore, we define a condition on $H^1(\mathcal{T})$ (and by extension, $H^1(\mathcal{T}, \mathbb{R}^d)$) called "asymptotic negligibility". Roughly speaking, a deformation is asymptotically negligible

if, up to an arbitrarily small error, it does not change the return vectors of large patches in a tiling. We then prove

Theorem 1.3 (Theorem 2.2)**.** *Let $f, g \in C^1(\Gamma, \mathbb{R}^d)$ be admissible shape functions for a fixed tiling space. If $\mathcal{I}(f) - \mathcal{I}(g)$ is asymptotically negligible, then there is a topological conjugacy between $\mathcal{T}_f$ and $\mathcal{T}_g$.*

Rescaling a tiling, or more generally applying a linear transformation, does not change the qualitative dynamical properties of the tiling space, such as mixing, minimality, or diffractivity. We therefore consider when a tiling space $\mathcal{T}_f$ is topologically conjugate to a linear transformation of a tiling space $\mathcal{T}_g$. Since linear transformations act naturally on the $\mathbb{R}^d$ factor of $H^1(\mathcal{T}, \mathbb{R}^d)$ we have:

Corollary 1.4. *The tiling spaces $\mathcal{T}_f$ and $\mathcal{T}_g$ are conjugate up to linear transformation if, for some linear transformation L, $\mathcal{I}(g) - L(\mathcal{I}(f))$ is asymptotically negligible.*

In Section 3 we specialize to nonperiodic tilings made from a primitive substitution, or a primitive substitution-with-amalgamation. The resulting spaces are orbit closures of self-similar or pseudo-self-similar tilings, respectively, in the terminology of [So1, So2, PS]. An essential object of study is the action of the substitution on the tiling space itself, and therefore on $H^1(\mathcal{T}, \mathbb{R}^d)$. We decompose $H^1(\mathcal{T}, \mathbb{R}^d)$ into (generalized) eigenspaces of the substitution operator. Let $S(\mathcal{T})$ be the span of the (generalized) eigenvectors with eigenvalue strictly less than one in magnitude. We show that the set of asymptotically negligible classes is precisely $S(\mathcal{T})$, and prove the following local converse to Theorem 1.3:

Theorem 1.5 (Theorem 3.3)**.** *For each shape parameter f there is a neighborhood U_f of $\mathcal{I}(f)$ such that, if $\mathcal{I}(g) \in U_f$ and $\mathcal{T}_f$ and $\mathcal{T}_g$ are topologically conjugate, then $\mathcal{I}(g) - \mathcal{I}(f) \in S(\mathcal{T})$.*

If there is a pseudo-self-similar tiling in $\mathcal{T}_f$, then $\mathcal{I}(f)$ is a Perron-Frobenius eigenvector of the substitution applied to $H^1(\mathcal{T}, \mathbb{R}^d)$. The Perron-Frobenius eigenspace, denoted $PF(\mathcal{T})$, is d^2 dimensional, and precisely equals $L\mathcal{I}(f)$, where L ranges over all linear transformations of $\mathbb{R}^d$. Corollary 1.4 has the following simple restatement for substitution tilings:

Corollary 1.6. *If $\mathcal{T}_f$ contains a pseudo-self-similar tiling, and if $\mathcal{I}(g) \in PF(\mathcal{T}) \oplus S(\mathcal{T})$, then $\mathcal{T}_g$ is topologically conjugate to a linear transformation applied to $\mathcal{T}_f$.*

In particular, if $H^1(\mathcal{T}, \mathbb{R}^d) = PF(\mathcal{T}) \oplus S(\mathcal{T})$, then all choices of shape give rise to conjugate dynamics, up to linear transformation. This case is the natural generalization, to higher dimensions, of Pisot substitutions.

For the space of Penrose tilings, $H^1(\mathcal{T}, \mathbb{R}^2)$ is 10-dimensional [AP]. The eigenvalues of the substitution are the golden mean τ, with multiplicity 4, $1 - \tau$ with multiplicity 4, and -1 with multiplicity 2. However, the -1 eigenspace is odd under rotation by 180 degrees, while the τ and $1 - \tau$ eigenspaces are even [ORS]. Any choice of shapes that preserves the (statistical) 2-fold rotational symmetry of the Penrose tiling must not involve the -1 eigenspace. As a result, the tiling space constructed from the tiling of

figure 2 is topologically conjugate to a linear transformation applied to the undistorted Penrose tiling space.

The techniques of Section 3 require recognizability, and hence draw heavily on Solomyak's generalization [So2] of Mossés work on 1-dimensional subshifts [M1, M2].

In Section 4 we study the spectra of substitution tilings and their deformations. We provide general criteria for the existence of point spectrum of translations acting on $\mathcal{T}_f$, similar in spirit to criteria found in [So1] for self-similar tilings and to criteria in [CS] for 1-dimensional substitution tiling spaces and their deformations. We also provide constraints on the form of that spectrum, in terms of the aforementioned decomposition of $\mathcal{I}(f)$ into eigenspaces of the substitution operator. In particular, we show:

Theorem 1.7 (Theorem 4.4). *If $H^1(\mathcal{T}, \mathbb{R}^d) \neq PF(\mathcal{T}) \oplus S(\mathcal{T})$, then for a generic choice of shape parameter f, $\mathcal{T}_f$ is topologically weakly mixing.*

Finally, in Section 5 we revisit the problem in one dimension, and recast the results of [CS] in topological terms.

2. The Map $\mathcal{I}$

To define the map $\mathcal{I}$ we first recall the inverse limit structure of tiling spaces, developed by Anderson and Putnam [AP] for substitution tilings, and generalized by Gähler [G] to apply to all translationally finite tilings. (See [BBG] for an alternate approach, and [ORS, Sa, BG] for further generalizations).

We first construct a complex Γ by taking the disjoint union of prototiles in the tiling, modulo identification of edges where tiles can meet. If somewhere in a tiling, edge i of a tile of type A is coincident with edge j of a tile of type B, then we identify edge i of A with edge j of B in the complex. (If $d > 2$ we also identify coincident faces and other structures of dimension up to $d-1$.)

We may also rewrite the tiling in terms of collared tiles, labeling each tile t by the patch consisting of all tiles that touch t. Applying the construction of the previous paragraph to the collared tiles gives a complex $\Gamma^{(1)}$. Collaring the collared tiles and applying the construction gives a complex $\Gamma^{(2)}$, and more generally applying the construction to k-times collared tiles gives a complex $\Gamma^{(k)}$.

There is a natural "forgetful" map α_0 from $\Gamma^{(1)}$ to Γ that simply ignores the collaring, and likewise a forgetful map $\alpha_k : \Gamma^{(k+1)} \mapsto \Gamma^{(k)}$. The inverse limit of the sequence of maps and spaces

$$\Gamma = \Gamma^{(0)} \xleftarrow{\alpha_0} \Gamma^{(1)} \xleftarrow{\alpha_1} \Gamma^{(2)} \cdots \tag{1}$$

is isomorphic to the tiling space $\mathcal{T}$ [G]. A point in Γ tells how to place a tile around the origin. A point in $\Gamma^{(1)}$ tells how to place a collared tile, i.e., a tile and its nearest neighbors. As k increases, the points in $\Gamma^{(k)}$ tell how to place larger and larger patches around the origin, and the entire sequence $(x_0, x_1, \ldots)$ with $x_i \in \Gamma^{(i)}$ and $x_i = \alpha_i(x_{i+1})$ tells how to place a complete tiling.

We have already seen that the a shape parameter f is a vector-valued 1-cochain in Γ. In fact, for any tile, $\delta f(t) = f(\partial t) = 0$, since this is the sum of the edge vectors around the tile t. Thus f is a cocycle, and defines a cohomology class $[f] \in H^1(\Gamma, \mathbb{R}^d)$.

Let $r_k : \mathcal{T} \mapsto \Gamma^{(k)}$ be the projection of the tiling space to the kth approximant. Define $\mathcal{I}(f) = r_0^*[f] \in H^1(\mathcal{T}, \mathbb{R}^d)$.

Theorem 2.1 (Theorem 1.2)**.** *Suppose f and g are admissible shape parameters with $\mathcal{I}(f) = \mathcal{I}(g)$. Then $\mathcal{T}_f$ and $\mathcal{T}_g$ are mutually locally derivable.*

Proof. We show how to construct a tiling in $\mathcal{T}_g$ from the corresponding tiling in $\mathcal{T}_f$ by a process that is completely local. The reverse process is of course similar.

Let $\pi_k = \alpha_0 \circ \alpha_1 \circ \cdots \circ \alpha_{k-1} : \Gamma^{(k)} \mapsto \Gamma$. If $\mathcal{I}(f) = \mathcal{I}(g)$, then for some finite k, $\pi_k^*[f] - \pi_k^*[g] = 0$, so $\pi_k^* f - \pi_k^* g = \delta\beta$ for some $\beta \in C^0(\Gamma^{(k)}, \mathbb{R}^d)$. Now each vertex $\mathbf{v}$ in a tiling in $\mathcal{T}_f$ maps to a unique vertex in $\Gamma^{(k)}$, the map being determined by a ball of size $(k+1)A$ around $\mathbf{v}$, where A is the diameter of the largest prototile. Moving each vertex $\mathbf{v}$ by $-\beta(\mathbf{v})$, and linearly interpolating the edges between vertices, converts a tiling in $\mathcal{T}_f$ to a tiling in $\mathcal{T}_g$. □

For a tiling $x \in \mathcal{T}$, a *recurrence* is an ordered pair $(\mathbf{z}_1, \mathbf{z}_2)$ of points in x such that $\mathbf{z}_1$ is a point in a tile and $\mathbf{z}_2$ is the corresponding point of a translate of that tile and such that there exist balls around $\mathbf{z}_1$ and $\mathbf{z}_2$ that agree (up to translation by $\mathbf{z}_2 - \mathbf{z}_1$, of course). If r is the supremum of the radii of the balls around $\mathbf{z}_1$ and $\mathbf{z}_2$ that agree, then we say the recurrence has *size* r. Each path along edges from $\mathbf{z}_1$ to $\mathbf{z}_2$ lifts to a closed loop in Γ, and hence to a closed chain in $C_1(\Gamma)$. Different paths from $\mathbf{z}_1$ to $\mathbf{z}_2$ correspond to homologous chains. The class in $H_1(\Gamma)$ of a recurrence is called a *recurrence class.* Note that recurrences of size greater than $(k+1)A$, where A is the diamater of the largest prototile, also lift to closed paths in $\Gamma^{(k)}$ and define classes in $H_1(\Gamma^{(k)})$. Since the tiling space $\mathcal{T}$ is assumed to be minimal, the set of recurrence classes is the same for every tiling in the space, and so we can speak of the recurrence classes of the tiling *space.*

An element η of $H^1(\Gamma^{(k)}, \mathbb{R})$, or of $H^1(\Gamma^{(k)}, \mathbb{R}^d)$, is said to be *asymptotically negligible* if, for each $\epsilon > 0$ there exists an constant R_ϵ such that η, applied to any recurrence of size greater than R_ϵ, is less than ϵ in magnitude. A class in $H^1(\mathcal{T}, \mathbb{R}^d)$ is asymptotically negligible if it is the pullback of an asymptotically negligible class in $H^1(\Gamma^{(k)}, \mathbb{R}^d)$ for some finite k. It is clear that any linear combination of asymptotically negligible classes is asymptotically negligible, so these classes form a subspace of $H^1(\mathcal{T}, \mathbb{R}^d)$, denoted $N(\mathcal{T})$.

Theorem 2.2 (Theorem 1.3)**.** *Let $f, g \in C^1(\Gamma, \mathbb{R}^d)$ be admissible shape functions for a fixed tiling space. If $\mathcal{I}(f) - \mathcal{I}(g)$ is asymptotically negligible, then there is a topological conjugacy $\phi : \mathcal{T}_f \mapsto \mathcal{T}_g$.*

Proof. We construct the conjugacy ϕ in stages. First pick a reference tiling $x \in \mathcal{T}_f$ that has a vertex at the origin.

For every vertex $\mathbf{v}$ in x there is a path $p_\mathbf{v}$ from the origin to $\mathbf{v}$ along edges, and the location of $\mathbf{v}$ is precisely $f(p_\mathbf{v})$. In $\phi(x)$ we place the corresponding vertex at $g(p_\mathbf{v})$. The path $p_\mathbf{v}$ is not uniquely defined, but different choices differ by boundaries, so the values of $f(p_\mathbf{v})$ and $g(p_\mathbf{v})$ are uniquely determined. Once the location of the vertices of $\phi(x)$ are specified, constructing the edges and tiles is straightforward.

This defines $\phi(x)$. For $\mathbf{z} \in \mathbb{R}^d$, let $\phi(x - \mathbf{z}) = \phi(x) - \mathbf{z}$. It remains to show that ϕ is uniformly continuous on the orbit of x, and hence can be extended to all of $\mathcal{T}$.

Suppose that $\mathbf{z}_1$ and $\mathbf{z}_2$ are vertices in x such that $x-\mathbf{z}_1$ and $x-\mathbf{z}_2$ agree on a large ball around the origin. Note that $\phi(x)-g(p_{\mathbf{z}_1})$ agrees exactly with $\phi(x)-g(p_{\mathbf{z}_2})$ on a large ball around the origin. Since $f-g$ is asymptotically negligible, $f(p_{\mathbf{z}_2})-f(p_{\mathbf{z}_1})$ is very close to $g(p_{\mathbf{z}_2})-g(p_{\mathbf{z}_1})$, so $\phi(x-\mathbf{z}_1)=\phi(x)-f(p_{\mathbf{z}_1})$ agrees with $\phi(x-\mathbf{z}_2)=\phi(x)-f(p_{\mathbf{z}_2})$ on a large ball around the origin, up to a small translation, of size $f(p_{\mathbf{z}_2})-f(p_{\mathbf{z}_1})-(g(p_{\mathbf{z}_2})-g(p_{\mathbf{z}_1}))$. Since this translation can be made arbitrarily small by making the recurrence of sufficiently large size (independent of the choice of $\mathbf{z}_1$ and $\mathbf{z}_2$), ϕ is uniformly continuous on the orbit of x.

To see that ϕ is invertible, construct a semi-conjugacy $\phi' : \mathcal{T}_g \mapsto \mathcal{T}_f$ by the same procedure, with the roles of f and g reversed, and with $\phi(x)$ as the reference tiling. Since $\phi'(\phi(x))=x$, it is clear that ϕ and ϕ' are inverses. $\square$

Note that $\mathcal{T}_f$ and $\mathcal{T}_g$ are typically not MLD. If two tilings in $\mathcal{T}_f$ agree on a large ball around the origin, the corresponding tilings in $\mathcal{T}_g$ agree on a large ball *up to a small translation*. Only if $f-g$ vanishes on large recurrence classes do they agree without translation. In that case, however, the following theorem shows that $\mathcal{I}(f)=\mathcal{I}(g)$, and we are back in the situation of Theorem 1.2.

Theorem 2.3. *Let $\beta \in H^1(\Gamma,\mathbb{R})$. If β vanishes on all recurrence classes of size greater than a fixed value R, then the pullback of β to $H^1(\mathcal{T},\mathbb{R})$ is zero.*

Proof. We use Kellendonk's P-equivariant cohomology [K], which relates the real-valued cohomology of a tiling space to closed and exact forms on a single tiling x, meeting some equivariance conditions. A differential form on the tiling x (viewed as a decorated copy of $\mathbb{R}^d$) is said to be P-equivariant if there is some radius r such that the value of the form at each point depends only on the tiling in a ball of radius r around that point. That is, if ρ is an equivariant form with radius r, and if $x-\mathbf{z}_1$ and $x-\mathbf{z}_2$ agree on a ball of radius r around the origin, then $\rho(\mathbf{z}_1)=\rho(\mathbf{z}_2)$. Kellendonk proves a version of the de Rham theorem: the closed P-equivariant k-forms on x, modulo d of the P-equivariant $(k-1)$-forms, is isomorphic to $H^k(\mathcal{T},\mathbb{R})$.

Now consider our class β, and represent it as a closed P-equivariant 1-form ρ. Let $\gamma(\mathbf{z})=\int_0^{\mathbf{z}}\rho$. It is clear that $\rho=d\gamma$, and the condition that β vanishes on large recurrence classes implies that γ is P-equivariant with radius R. Thus ρ represents the zero class in $H^1(\mathcal{T},\mathbb{R})$. $\square$

3. Substitution Tiling Spaces

We now specialize to substitution tiling spaces, where we can obtain sharper results. In this section we identify the asymptotically negligible cohomology classes and prove Theorem 1.5.

A substitution tiling system is determined by a substitution σ from the set of prototiles to the set of finite patches, such that for each prototile t, the tiles of $\sigma(t)$ do not overlap, and such that their union is the rescaled prototile λt, where λ is a fixed "stretching factor". For example, in the "chair" tiling, there are four prototiles: One is the L-shaped tile of figure 3 and the others are this tile rotated by $\pi/2$, π and $3\pi/2$. The substitution map is shown in figure 3. We extend the map σ to patches, dilating the entire patch by

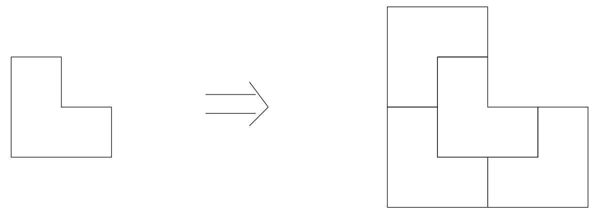

Figure 3. The chair tile and its substitution

a factor of λ and replacing each dilated tile λt_i with $\sigma(t_i)$. For each prototile t we have a sequence of patches $t, \sigma(t), \sigma^2(t), \cdots$. A tiling x of $\mathbb{R}^d$ by prototiles is called admissible is for every finite patch P of x, there exists a prototile t and an integer n such that P is a translate of a subpatch of $\sigma^n(t)$. The substitution tiling space $\mathcal{T}$ is then the set of all admissible tilings, and σ extends to a continuous map $\mathcal{T} \mapsto \mathcal{T}$.

We assume that the substitution is primitive. That is, there exists an integer n such that, for any two (possibly identical) prototiles t_1 and t_2, $\sigma^n(t_1)$ contains a copy of t_2. This ensures that $\mathcal{T}$ is minimal. Furthermore, we assume that $\mathcal{T}$ is non-periodic. There does not exist a tiling $x \in \mathcal{T}$ and a nonzero $\mathbf{z} \in \mathbb{R}^d$ with $x - \mathbf{z} = x$. By a theorem of Solomyak [So2], extending previous work by Mossé [M1, M2], this implies that $\sigma : \mathcal{T} \mapsto \mathcal{T}$ is a homeomorphism. There then exists a "recognition length" D such that, if x and y are tilings that agree on a ball of radius $r > D$ about a point $\mathbf{z} \in \mathbb{R}^d$, then $\sigma^{-1}(x)$ and $\sigma^{-1}(y)$ agree on the tile(s) containing the point $\mathbf{z}/\lambda$.

In every substitution tiling space there exists a fixed point of some power of the substitution. Replacing σ with a power of σ does not change the tiling space, so we can assume, without loss of generality, that there is a tiling x with $\sigma(x) = x$. Such a tiling is called "self-similar", and has the property that λx is a tiling by large tiles, each of which is a union of tiles of x. Some authors begin by defining self-similar tilings by this property, and then take $\mathcal{T}$ to be the closure of the translational orbit of x.

Closely related to self-similar tilings are pseudo-self-similar tilings. A tiling x is pseudo-self-similar if, for some scaling factor λ, λx and x are MLD. Natalie Priebe Frank and Boris Solomyak [PS] have shown that every pseudo-self-similar planar tiling with polygonal tiles is MLD to a self-similar tiling. However, the tiles of the self-similar tiling may not be polygonal; rather, they may have fractile boundaries. Conversely, Natalie Priebe Frank [Pr] showed (in arbitrary dimension) that each self-similar tiling (with tiles of arbitrary shape) is MLD to a pseudo-self-similar tiling with polyhedral tiles meeting full-face to full-face.

Orbit closures of pseudo-self-similar tilings may also be viewed as coming from a substitution-with-amalgamation. This is a prescription for replacing each rescaled prototile λt by a collection of tiles, such that the images of the different tiles in a tiling do not overlap and do not leave any gaps. As with an ordinary substitution, σ extends to a homeomorphism $\mathcal{T} \mapsto \mathcal{T}$. The only difference is that the tiles of $\sigma(x)$ may stick in and out of the tiles of λx. An example of a substitution-with-amalgamation, for a tiling by marked hexagonal tiles, is shown in figure 4.

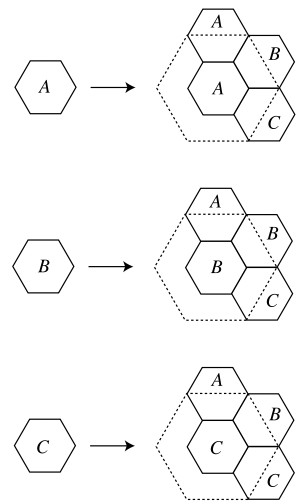

FIGURE 4. A substitution with amalgamation

In this section we assume that $\mathcal{T}$ is a nonperiodic substitution tiling space derived from a substitution-with-amalgamation σ. Since σ maps $\mathcal{T}$ to $\mathcal{T}$, it induces a pullback map: $\sigma^* : H^1(\mathcal{T}, \mathbb{R}^d) \mapsto H^1(\mathcal{T}, \mathbb{R}^d)$. We decompose $H^1(\mathcal{T}, \mathbb{R}^d)$ into (generalized) eigenspaces of σ^*. The largest eigenvalue, denoted λ_{PF}, equals the stretching factor and has multiplicity d^2. (In terms of P-equivariant cohomology, a basis for this space is $dz^i \otimes e_j$, where $e_1, \ldots, e_d$ is the standard basis for $\mathbb{R}^d$ and $z^1, \ldots z^d$ are Cartesian coordinates.) The corresponding eigenspace is denoted $PF(\mathcal{T})$. The span of the (generalized) eigenvectors with eigenvalues of magnitude strictly less than 1 is denoted $S(\mathcal{T})$.

Anderson and Putnam's results [AP] were originally stated for substitutions without amalgamation, but apply equally well to substitutions with amalgamation. The substitution σ induces a map from $\Gamma^{(1)}$ to itself, and $\mathcal{T}$ is the inverse limit of the sequence of maps:

$$\Gamma^{(1)} \xleftarrow{\sigma} \Gamma^{(1)} \xleftarrow{\sigma} \Gamma^{(1)} \xleftarrow{\sigma} \Gamma^{(1)} \cdots \tag{2}$$

The cohomology $H^1(\mathcal{T}, \mathbb{R}^d)$ is then the direct limit of $H^1(\Gamma^{(1)}, \mathbb{R}^d)$ under $\sigma^* : H^1(\Gamma^{(1)}, \mathbb{R}^d) \mapsto H^1(\Gamma^{(1)}, \mathbb{R}^d)$. Nonzero eigenspaces of σ^* applied to $H^1(\mathcal{T}, \mathbb{R}^d)$ correspond to nonzero eigenspaces of σ^* applied to $H^1(\Gamma^{(1)}, \mathbb{R}^d)$. The only difference between $H^1(\mathcal{T}, \mathbb{R}^d)$ and $H^1(\Gamma^{(1)}, \mathbb{R}^d)$ is that $H^1(\Gamma^{(1)}, \mathbb{R}^d)$ may contain a (generalized) zero-eigenspace of σ^*, while $H^1(\mathcal{T}, \mathbb{R}^d)$ does not.

This decomposition into eigenspaces of σ^* makes it easy to identify the asymptotically negligible classes.

Theorem 3.1. *If $\mathcal{T}$ is a substitution tiling space, then $N(\mathcal{T}) = S(\mathcal{T})$.*

Proof. For simplicity, suppose that the action of σ^* on $H^1(\mathcal{T},\mathbb{R}^d)$ is diagonalizable. First we show that $N(\mathcal{T}) \subset S(\mathcal{T})$. For $\beta \in N(\mathcal{T})$, decompose β as

$$\beta = \beta_1 + \beta_2 + \cdots + \beta_k, \tag{3}$$

with all terms nonzero, and with each β_j an eigenvector of σ^* with eigenvalue λ_j, with $|\lambda_1| \geq |\lambda_2| \geq \cdots \geq |\lambda_k|$. Since $\beta_1 \neq 0$, by Theorem 2.3 there exists a recurrence class $[p_z]$ such that $\beta_1[p_z] \neq 0$. The size of the recurrence $\sigma^n(p_z)$ grows exponentially with n, so $\beta(\sigma^n(p_z))$ goes to zero as $n \to \infty$. However,

$$\beta(\sigma^n(p_z)) = ((\sigma^*)^n\beta)p_z = \lambda_1^n\beta_1(p_z) + \lambda_2^n\beta_2(p_z) + \cdots + \lambda_k^n\beta_k(p_z). \tag{4}$$

Since $\beta_1(p_z) \neq 0$, the only way this can converge to zero is if $|\lambda_1| < 1$, hence $\beta \in S(\mathcal{T})$.

To show that $S(\mathcal{T}) \subset N(\mathcal{T})$, note that every recurrence class $[p_z]$ can be written as a sum

$$[p_z] = \sum_{j=0}^{\infty} \sigma^j[p_j], \tag{5}$$

with only a finite number of nonzero terms, and with all the $[p_j]$'s belonging to a fixed bounded subset U of $H_1(\Gamma^{(1)})$. (The argument is essentially that found in [So2] and is not repeated here.) If the recurrence has size greater than $D(1+\lambda_{PF}+\cdots+\lambda_{PF}^{n-1})$ where D is the recognition length, then the first n terms are zero.

Suppose that β is an eigenvector of σ^* with eigenvalue λ of magnitude strictly less than 1. Let K be larger than the greatest value that β takes on U. Then for any recurrence class $[p_z]$, $|\beta(p_z)| < K/(1-|\lambda|)$. If p_z has size greater than $D(1+\lambda_{PF}+\cdots+\lambda_{PF}^{n-1})$, then $|\beta(p_z)| < K|\lambda|^n/(1-|\lambda|)$. Thus β is asymptotically negligible. Since linear combinations of asymptotically negligible classes are asymptotically negligible, $S(\mathcal{T}) \subset N(\mathcal{T})$.

When σ^* is not diagonalizable, the analysis is only slightly more complicated. The expansion of $\beta(\sigma^n(p_j))$ may involve n-th powers of eigenvalues times polynomials in n, rather than just n-th powers of eigenvalues. However, the conclusions are unchanged. □

Theorem 3.2. *If, for some non-negative integer k, either $\mathcal{I}(g) - (\sigma^*)^k\mathcal{I}(f) \in S(\mathcal{T})$ or $\mathcal{I}(f) - (\sigma^*)^k\mathcal{I}(g) \in S(\mathcal{T})$, then $\mathcal{T}_f$ is topologically conjugate to $\mathcal{T}_g$.*

Proof. The case where $k = 0$ was already proven by Theorems 1.3 and 3.1. All that remains is to prove conjugacy when $\mathcal{I}(g) = \sigma^*(\mathcal{I}(f))$. Since the MLD class of $\mathcal{T}_g$ depends only on $\mathcal{I}(g)$, we can assume that $g = \sigma^*(f)$ as cochains on $\Gamma^{(1)}$. [Note: since we need to work with $\Gamma^{(1)}$ rather than $\Gamma^{(0)}$, we pull our shape parameters f and g from Γ back to $\Gamma^{(1)}$, and henceforth view them as closed cochains in $C^1(\Gamma^{(1)},\mathbb{R}^d)$.] That is, a cluster of tiles $\sigma(t_i)$ in $\mathcal{T}_g$ has the same dispacements as the single tile t_i in $\mathcal{T}_f$. To convert a tiling $x \in \mathcal{T}_f$ to a tiling $\phi(x) \in \mathcal{T}_g$, replace each tile $t \in x$ by the cluster $\sigma(t)$, with sizes given by g. □

Theorem 3.2 gives sufficient conditions for two tiling spaces to be conjugate. For infinitesimal deformations, these conditions are also necessary:

Theorem 3.3 (Theorem 1.5)**.** *For each shape parameter f there is a neighborhood U_f of $\mathcal{I}(f)$ such that, if $\mathcal{I}(g) \in U_f$ and $\mathcal{T}_f$ and $\mathcal{T}_g$ are topologically conjugate, then $\mathcal{I}(g) - \mathcal{I}(f) \in S(\mathcal{T})$.*

The remainder of this section is devoted to proving Theorem 3.3. To prove the theorem we must understand the extent to which patterns in $\mathcal{T}$ can repeat themselves. For each recurrence $(\mathbf{z}_1, \mathbf{z}_2)$, let the *degree* of the recurrence be its size divided by $|\mathbf{z}_2 - \mathbf{z}_1|$. For fixed $p > 0$, we show that the recurrence classes corresponding to recurrences with degree p or greater can be grouped into a finite number of families, and we construct a conjugacy invariant from the asymptotic displacements of these recurrences. Comparing these invariants for different shape parameters then gives necessary conditions for conjugacy.

It is convenient to work with explicit matrices. The set of integer linear combinations of recurrence classes of size at least $2A$ (where A is the diameter of the largest tile) is a sub-lattice of $H_1(\Gamma^{(1)})$. Let $\{a_1, \ldots, a_s\}$ be a basis for this sub-lattice. Since σ maps recurrences to recurrences, it maps the lattice to itself. Let M be the matrix of this map relative to the basis $\{a_1, \ldots, a_s\}$. Note that the entries of M are integers. The matrix M will play a role similar to that of the substitution matrix in one-dimensional substitutions. For each recurrence class in $H_1(\Gamma^{(1)})$, the corresponding *recurrence vector* in $\mathbb{Z}^s$ is the decomposition of the class in the $\{a_1, \ldots, a_s\}$ basis. For any shape parameter f, let $\mathbf{L}_f = (f(a_1), \ldots, f(a_s))$ be the ($\mathbb{R}^d$-valued row) vector that gives the displacements, in $\mathbb{R}^d$, corresponding to lifts of the various loops. For a recurrence with vector $\mathbf{v}$, the corresponding displacement is $\mathbf{L}_f\mathbf{v}$. The Euclidean length of $\mathbf{L}_f\mathbf{v}$ is called the *length* of $\mathbf{v}$, and denoted $|\mathbf{v}|_f$. We call $\mathbf{L}_f$ the *shape vector* of $\mathcal{T}_f$.

Let $\mathcal{T}_0$ be a tiling space whose shape vector $\mathbf{L}_0$ is a left Perron-Frobenius eigenvector of M. This can be the original substitution tiling space, a linear transformation applied to the original space, or a space that is MLD to such a linear transformation. In all such cases, there exists a pseudo-self-similar tiling in the space, and we call the space itself pseudo-self-similar.

If $\mathbf{v}$ is a recurrence vector of degree p in $\mathcal{T}_0$, corresponding to a recurrence $(\mathbf{z}_1, \mathbf{z}_2)$ in a tiling x, then $M\mathbf{v}$ is also a recurrence vector of degree at least p, corresponding to the recurrence $(\lambda_{PF}\mathbf{z}_1, \lambda_{PF}\mathbf{z}_2)$ in the tiling $\sigma(x)$. Both the size of the matching balls around $\mathbf{z}_1$ and $\mathbf{z}_2$ and the distance from $\mathbf{z}_1$ to $\mathbf{z}_2$ get stretched by the same factor λ_{PF} in the substitution. Thus each recurrence vector of degree p gives rise to a family of recurrence vectors $M^k\mathbf{v}$. The following theorem limits the number of such families.

Theorem 3.4. *Let $\mathcal{T}_0$ be any pseudo-self-similar tiling space and let $p > 0$. There is a finite collection of vectors $\{\mathbf{v}_1, \ldots, \mathbf{v}_N\}$ such that every recurrence vector of degree at least p for $\mathcal{T}_0$ is of the form $M^k\mathbf{v}_i$ for some pair (k, i).*

Proof. As before, let D be the recognition length of the substitution σ acting on $\mathcal{T}_0$. Suppose $\mathbf{v}$ is a recurrence vector of degree at least p for $\mathcal{T}_0$, where the radius R of the matching balls is much greater than D. Then the supertiles within balls of radius $R - D$ around $\mathbf{z}_1$ and $\mathbf{z}_2$ must also agree. Thus there is a recurrence vector $\mathbf{v}_1$ such that $\mathbf{v} = M\mathbf{v}_1$, and such that $\mathbf{v}_1$ is a recurrence vector of degree at least $(R-D)/(\lambda_{PF}|\mathbf{v}_1|_0) \geq (p|\mathbf{v}|_0 - D)/(\lambda_{PF}|\mathbf{v}_1|_0) = p - (D/|\mathbf{v}|_0)$.

Repeating the process, we find recurrence vectors $\mathbf{v}_i$ such that $M\mathbf{v}_i = \mathbf{v}_{i-1}$ and $\mathbf{v}_i$ is a recurrence vector of degree at least

$$(6) \qquad p - \frac{D}{|\mathbf{v}_i|_0}(\lambda_{PF}^{-1} + \lambda_{PF}^{-2} + \cdots + \lambda_{PF}^{-i}) \geq p - \frac{D}{|\mathbf{v}_i|_0}\sum_{\ell=1}^{\infty}\lambda_{PF}^{-\ell} = p - \frac{D}{|\mathbf{v}_i|_0(\lambda_{PF}-1)}.$$

Now pick $\epsilon < p$. We have shown that every recurrence vector of degree at least p is of the form $M^k\mathbf{v}_i$, where $\mathbf{v}_i$ is a recurrence vector of degree at least $p-\epsilon$, and where $|\mathbf{v}_i|_0$ is bounded by $\frac{D\lambda_{PF}}{\epsilon(\lambda_{PF}-1)}$. However, by finite local complexity, there are only a finite number of recurrence vectors of this length. □

As defined, the degree of a recurrence vector depends on the shape vector of the tiling space. Indeed, since linear transformations do not preserve lengths or ratios of lengths, the degree of a recurrence vector in two pseudo-self-similar tiling spaces may not be the same. However, for any tiling space $\mathcal{T}_f$, the degrees of large recurrence vectors are approximately the same as that of some pseudo-self-similar tiling space:

Lemma 3.5. *Let $\mathbf{L}_f$ be decomposed as $\mathbf{L} = \mathbf{L}_0 + \mathbf{L}_r$, where $\mathbf{L}_r$ is a linear combination of (generalized) eigenvectors with eigenvalue less than λ_{PF}, and $\mathbf{L}_0$ is a Perron-Frobenius eigenvector. Let $\mathcal{T}_0$ be a pseudo-self-similar tiling space with shape vector $\mathbf{L}_0$. For each recurrence vector $\mathbf{v}$, let $r_{\mathbf{v}} = |\mathbf{v}|_f/|\mathbf{v}|_0$. For each $\epsilon > 0$ there exists a length R such that all recurrences vectors v with $|\mathbf{v}|_f > R$ have $|r_{\mathbf{v}} - 1| < \epsilon$.*

Recurrences of the form $M^n\mathbf{v}$, with n large, have lengths in $\mathcal{T}_f$ and $\mathcal{T}_0$ that differ only slightly. Since any long recurrence is the sum of high order recurrences plus a bounded number of lower order recurrences, the result follows. □

Lemma 3.6. *Let $\epsilon > 0$, and let f be fixed. There is a self-similar tiling space $\mathcal{T}_0$ and a length R such that every recurrence vector of degree at least p in $\mathcal{T}_f$, of length greater than R, is a recurrence vector of degree at least $p-\epsilon$ in $\mathcal{T}_0$. Furthermore, every recurrence vector of degree at least $p+\epsilon$ in $\mathcal{T}_0$, of length greater than R, is a recurrence vector of degree at least p in $\mathcal{T}_f$.*

Proof. Decompose $\mathbf{L}_f = \mathbf{L}_0 + \mathbf{L}_r$ as above, and let $\mathcal{T}_0$ be the pseudo-self-similar space with length vector $\mathbf{L}_0$. By Lemma 3.5, the ratio of lengths in $\mathcal{T}_f$ and $\mathcal{T}_0$ approaches 1 for large recurrences. Thus the ratio of the radii of the balls around $\mathbf{z}_1$ and $\mathbf{z}_2$ and the distance $|\mathbf{z}_2 - \mathbf{z}_1|$ are within ϵ for large recurrences. □

Theorem 3.4 showed there there are only a finite number of families of recurrence vectors of a given degree for a self-similar tiling. Lemma 3.6 extends that result to all tiling spaces.

We now construct a topological invariant from the asymptotic displacements $\mathbf{L}_f\mathbf{v}$ of large recurrence vectors $\mathbf{v}$.

Theorem 3.7. *Suppose that $\phi : \mathcal{T}_f \mapsto \mathcal{T}_g$ is a topological conjugacy. Given any positive constants $p, \epsilon_1, \epsilon_2$, there exists a positive constant R such that, for each recurrence vector*

$\mathbf{v}$ *of degree at least* p *for* $\mathcal{T}_g$ *with* $|\mathbf{v}|_f > R$*, there exists a recurrence vector* $\mathbf{v}'$ *of degree at least* $p - \epsilon_1$ *for* $\mathcal{T}_g$ *with* $|\mathbf{L}_f\mathbf{v} - \mathbf{L}_g\mathbf{v}'| < \epsilon_2$*.*

In other words, up to small errors in the degree and the displacement that vanish in the limit of large recurrence classes, the degrees and displacements of the recurrence classes are conjugacy invariants.

Proof. Let A be the diameter of the largest tile in the $\mathcal{T}_g$ system. Since ϕ is uniformly continuous (being continuous with a compact domain), there exists a constant D_0 such that, if x and y are tilings in $\mathcal{T}_f$ that agree on a ball of radius D_0 around the origin, then $\phi(x)$ and $\phi(y)$ agree on a ball or radius A around the origin, up to a small translation. Thus if x and y agree on a ball of radius $R > D_0$, then $\phi(x)$ and $\phi(y)$ agree on a ball of radius $R - D_0$, up to a translation whose norm is bounded by a decreasing function $h(R)$, with $\lim_{R\to\infty} h(R) = 0$.

Now let $\mathbf{v}$ be a recurrence vector of degree p for $\mathcal{T}_f$, representing a recurrence $(\mathbf{z}_1, \mathbf{z}_2)$ in a tiling $x \in \mathcal{T}_f$. Then $x - \mathbf{z}_1$ and $x - \mathbf{z}_2$ agree on a ball of radius $p|\mathbf{v}|_f$ about the origin, so $\phi(x) - \mathbf{z}_1$ and $\phi(x) - \mathbf{z}_2$ agree on a ball of radius $p|\mathbf{v}|_f - D_0$, up to translation of size at most $h(R)$. Thus there exists a point $\mathbf{z}_3$, within $h(D)$ of $\mathbf{z}_2$, such that $\mathbf{v}' = [(\mathbf{z}_1, \mathbf{z}_3)]$ is a recurrence class of degree at least $\frac{(p|\mathbf{v}|_f - D_0}{|\mathbf{v}|_f + h(D)}$ in $\mathcal{T}_g$. For D large enough, this is greater than $p - \epsilon_1$ and $|\mathbf{L}_f\mathbf{v} - \mathbf{L}_g\mathbf{v}'| < h(R)$ is less than ϵ_2. $\square$

Proof of Theorem 3.3.

Note that all recurrence vectors are recurrence vectors of some positive degree. We can therefore pick p_0 such that the integer span of the recurrence vectors of degree at least p_0 is an s-dimensional sublattice of $\mathbb{Z}^s$. Pick ϵ small enough (and adjust p_0 by up to ϵ, if necessary) so that all the families of recurrence vectors of degree at least $p_0 - \epsilon$ are also families of recurrence vectors of degree at least $p_0 + \epsilon$. Let $\mathbf{v}_1, \ldots, \mathbf{v}_k$ be generating vectors of those families. By multiplying by appropriate powers of M, these can be chosen so that all of the magnitudes of the displacements $|M^k\mathbf{v}_i|_f$ are within a factor of λ_{PF} of one another for large k.

Now pick a neighborhood U_ϵ of f in $C^1(\Gamma, \mathbb{R}^d)$ small enough that, if $g' \in U_\epsilon$, then every recurrence class in $\mathcal{T}_f$ of degree at least $p_0 + \epsilon$ is also a recurrence class in $\mathcal{T}_{g'}$ of degree at least p_0, and such that every recurrence class in $\mathcal{T}_{g'}$ of degree at least p_0 is a recurrence class in $\mathcal{T}_f$ of degree at least $p_0 - \epsilon$. This insures that the families of recurrence classes of degree at least p_0 in the two tiling spaces are exactly the same.

If $\mathcal{T}_f$ and $\mathcal{T}_{g'}$ are conjugate then, by Theorem 3.7, the displacements $\mathbf{L}_f M^k \mathbf{v}_j$ can be approximated by $\mathbf{L}_{g'} M^{k'} \mathbf{v}_{j'}$ for some k', j', and the approximation get successively better as $k \to \infty$. However, if U_ϵ is chosen small enough, the only values of k', j' that come close to approximating are $k' = k$ and $j' = j$. Thus the limit of $(\mathbf{L}_f - \mathbf{L}_{g'})M^k\mathbf{v}_j$ must be zero. By Theorem 2.3, this implies that $(\sigma^*)^k(\mathcal{I}(f) - \mathcal{I}(g'))$ approaches zero as $k \to \infty$, and hence that $\mathcal{I}(f) - \mathcal{I}(g') \in S(\mathcal{T})$.

Finally, let $U_f = \mathcal{I}(U_\epsilon)$. If $\mathcal{I}(g) \in U_f$, then $\mathcal{T}_g$ is MLD to a tiling space $\mathcal{T}_{g'}$ with $g' \in U_\epsilon$ and $\mathcal{I}(g') = \mathcal{I}(g)$. Since $\mathcal{T}_{g'}$ is conjugate to $\mathcal{T}_f$, $\mathcal{I}(f) - \mathcal{I}(g) = \mathcal{I}(f) - \mathcal{I}(g') \in S(\mathcal{T})$. $\square$

4. Ergodic Properties

Now we turn to the topological point spectrum of substitution tiling spaces, by which we mean the eigenvalues of continuous eigenfuctions of the translation action. First we determine general criteria for a vector to be an eigenvalue for a continuous eigenfunction, and then we apply this criteria to some special cases, depending on the form of the matrix M defined in Section 3. This eventually leads to criteria for topological weak mixing.

Theorem 4.1. *The vector* $\mathbf{k} \in \mathbb{R}^d$ *is in the point spectrum of* $\mathcal{T}_f$ *if and only if, for every recurrence vector* $\mathbf{v}$*,*

$$\frac{1}{2\pi}(\mathbf{k} \cdot \mathbf{L}_f) M^m \mathbf{v} \to 0 \textit{(mod 1) as } m \to \infty, \tag{7}$$

where the convergence is uniform in the size of $\mathbf{v}$*.*

Proof. Let x_0 be a tiling in $\mathcal{T}_0$ fixed by the substitution homeomorphism, let x be its image in $\mathcal{T}_f$ under the homeomorphism of Theorem 1.1, and let $E : \mathcal{T}_f \mapsto S^1$ be a continuous eigenfunction with eigenvalue $\mathbf{k}$. Let $\mathbf{v}$ be a recurrence vector, then there is a recurrence $(\mathbf{z}_1, \mathbf{z}_2)$ in x with vector $\mathbf{v}$ of some size s_0. By applying the substitution homeomorphism m times, we obtain a recurrence $(\mathbf{z}_1^m, \mathbf{z}_2^m)$ in x with recurrence vector $M^m\mathbf{v}$, whose displacement is $\mathbf{z}_2^m - \mathbf{z}_1^m = \mathbf{L}_f M^m \mathbf{v}$. Then $x - \mathbf{z}_1^m$ and $x - \mathbf{z}_2^m$ agree on patches of size s_m, where $s_m \to \infty$ as $m \to \infty$. Hence

$$1 = \lim_{m\to\infty} \frac{E(x - \mathbf{z}_2^m)}{E(x - \mathbf{z}_1^m)} = \lim_{m\to\infty} \frac{E(x)\exp(-i\mathbf{k}\cdot\mathbf{z}_2^m)}{E(x)\exp(-i\mathbf{k}\cdot\mathbf{z}_1^m)} = \lim_{m\to\infty} \exp(-i\mathbf{k}\cdot(\mathbf{z}_2^m - \mathbf{z}_1^m)). \tag{8}$$

Thus, we obtain Equation 7, and the uniform convergence follows from the uniform continuity of E.

Conversely, assume that we have the stated convergence for all recurrence vectors $\mathbf{v}$ for some $\mathbf{k} \in \mathbb{R}^d$. We construct a continuous eigenfunction E by first assigning x the value 1. Then we necessarily have for any $\mathbf{z} \in \mathbb{R}^d$, $E(x - \mathbf{z}) = \exp(-i\mathbf{k}\cdot\mathbf{z})$. To show that E extends as required to all of $\mathcal{T}_f$, it suffices to show that E as so defined is uniformly continuous on the orbit of x. But, given an $\epsilon > 0$, by Equation 7 if $x - \mathbf{z}_1$ and $x - \mathbf{z}_2$ agree on patches of sufficiently large size up to a small translation, we have that $E(x - \mathbf{z}_1)$ and $E(x - \mathbf{z}_2)$ agree to within ϵ. □

The application of this criterion depends on the eigenvalues of M and on the possible forms of the recurrence vectors.

Theorem 4.2. *Suppose that all the eigenvalues of* M *are of magnitude 1 or greater. If* $\mathbf{k}$ *is in the spectrum, then all elements of* $\mathbf{k} \cdot \mathbf{L}_f / 2\pi$ *are rational.*

Proof. Let $\mathbf{k}$ be in the point spectrum, and consider the sequence of real numbers $t_m = (\mathbf{k}\cdot\mathbf{L}_f) M^m \mathbf{v}/(2\pi)$, where $\mathbf{v}$ is a fixed recurrence vector. Let $p(\lambda) = \lambda^n + a_{n-1}\lambda^{n-1} + \cdots + a_0$

be the characteristic polynomial of M. Note that the a_i's are all integers, since M is an integer matrix. Since $p(M) = 0$, the t_m's satisfy the recursion:

$$t_{m+n} = -\sum_{k=0}^{n-1} a_k t_{m+k}. \tag{9}$$

By Theorem 4.1, the t_m's converge to zero $\pmod 1$. That is, we can write

$$t_m = i_m + r_m \tag{10}$$

where the i_m's are integers, and the r_m's converge to zero as real numbers. By substituting the division (10) into the recursion (9), we see that both the i's and the r's must separately satisfy the recursion (9), once m is sufficiently large. However, any solution to this recursion relation is a linear combination of powers of the eigenvalues of M (or polynomials in m times eigenvalues to the m-th power, if M is not diagonalizable). Since the eigenvalues are all of magnitude one or greater, such a linear combination converges to zero only if it is identically zero. Therefore r_m must be identically zero for all sufficiently large values of m.

Apply this procedure to s linearly independent recurrence vectors $\mathbf{v}_1, \ldots, \mathbf{v}_s$, and pick m large enough that the corresponding $t_m(\mathbf{v}_i)$ are integers for each $i = 1, \ldots, s$. Note that t_m is an integer linear combination of the elements of the vector $\mathbf{k} \cdot \mathbf{L}_f/(2\pi)$. However,

$$(t_m(\mathbf{v}_1), \ldots, t_m(\mathbf{v}_s)) = (\mathbf{k} \cdot \mathbf{L}_f/2\pi) M^m (\mathbf{v}_1, \ldots, \mathbf{v}_s) \tag{11}$$

The matrices M and $V = (\mathbf{v}_1, \ldots, \mathbf{v}_s)$ are invertible and have integer entries, so by Cramer's rule their inverses have rational entries. Thus the components of $\mathbf{k} \cdot \mathbf{L}_f/(2\pi)$ must all be rational. □

Corollary 4.3. *Suppose all the eigenvalues of M have magnitude 1 or greater. Let $G \subset \mathbb{R}^d$ be the free Abelian group generated by the entries of $\mathbf{L}_f$. Let $\bar{G}$ be the closure of G in $\mathbb{R}^d$, and let P be the identity component of $\bar{G}$. The spectrum of $\mathcal{T}_f$ lies in the orthogonal complement of P.*

In dimension greater than 1, the assumptions of Theorem 4.2 are rarely met. When eigenvalues of magnitude less than 1 exist, the conclusions are somewhat weaker.

Theorem 4.4. *Let S be the span of the (generalized left-) eigenspaces of M with eigenvalues of magnitude strictly less than 1. If $\mathbf{k}$ is in the point spectrum, then $\mathbf{k} \cdot \mathbf{L}_f/2\pi$ is the sum of a rational vector and an element of S.*

Proof. Let V be the matrix $(\mathbf{v}_1, \ldots, \mathbf{v}_s)$ as in the proof of Theorem 4.2. Construct the row vector

$$\mathbf{t}_m = \frac{1}{2\pi}(\mathbf{k} \cdot \mathbf{L}_f) M^m V = \frac{1}{2\pi}(\mathbf{k} \cdot \mathbf{L}_f) V (V^{-1} M V)^m. \tag{12}$$

As in the proof of Theorem 4.2, each entry of $\mathbf{t}_m$ converges to zero (mod 1), so we can write $\mathbf{t}_m = \mathbf{i}_m + \mathbf{r}_m$, with each entry of $\mathbf{i}_m$ integral and $\mathbf{r}_m$ converging to zero, and with the eventual conditions

$$\mathbf{i}_{m+1} = \mathbf{i}_m V^{-1} M V; \quad \mathbf{r}_{m+1} = \mathbf{r}_m V^{-1} M V. \tag{13}$$

Since the $\mathbf{r}_m$'s converge to zero, they must lie in the span of the small eigenvalues of $V^{-1}MV$. Thus, by adding an element of S to $(\mathbf{k}\cdot\mathbf{L}_f)/2\pi$, we can then get all the $\mathbf{r}_m$'s to be identically zero beyond a certain point. Since M is invertible, the resulting value of $\mathbf{k}\cdot\mathbf{L}_f/2\pi$ must then be rational. □

Another way of stating the same result is to say that $(\mathbf{k}\cdot\mathbf{L}_f)/2\pi$, projected onto the span of the large eigenvectors, equals the projection of a rational vector onto this span.

This theorem can be used in two different ways. First, it constrains the set of shape parameters (that is, vectors $\mathbf{L}_f$) for which the system admits point spectrum. Let d_b be the number of large eigenvalues, counted with (algebraic) multiplicity. There are only a countable number of possible values for the projection of $\mathbf{k}\cdot\mathbf{L}_f/2\pi$ onto the span of the large (generalized) eigenvectors. In other words, one must tune d_b parameters to a countable number of possible values in order to achieve a point in the spectrum. Of course, $\mathbf{k}$ itself gives d parameters. Thus we must tune at least $d_b - d$ additional parameters to have any point spectrum at all. In particular, if $d_b > d$, then a generic choice of shape parameter gives topological weak mixing, proving Theorem 1.7. (Note that the Perron-Frobenius eigenvector always occurs with multiplicity d, so that $d_b - d$ is never negative.)

A second usage is to constrain the spectrum for fixed $\mathbf{L}_f$. The rational points in $\mathbb{R}^s$, projected onto the span of the large eigenvalues, and then intersected with the d-plane defined by a fixed $\mathbf{L}_f$ (i.e., the set of all possible products $(\mathbf{k}\cdot\mathbf{L}_f)/2\pi$), forms a vector space over $\mathbb{Q}$ of dimension at most $s+d-d_b$. As a result, the point spectrum tensored with $\mathbb{Q}$ is a vector space over $\mathbb{Q}$ whose dimension is bounded by d plus the number of small eigenvalues. Below we derive an even stronger result, in which only the small eigenvalues that are conjugate to the Perron-Frobenius eigenvalue contribute to the complexity of the spectrum.

Theorem 4.5. *Let b_{PF} be the number of large eigenvalues, counted without multiplicity, that are algebraically conjugate to the Perron-Frobenius eigenvalue λ_{PF} (including λ_{PF} itself), and let s_{PF} be the number of small eigenvalues conjugate to λ_{PF}. For fixed $\mathbf{L}_f$, the dimension over $\mathbb{Q}$ of the point spectrum tensored with $\mathbb{Q}$ is at most $d(s_{PF}+1)$.*

Proof. As a first step we diagonalize M over the rationals as far as possible. By rational operations we can always put M in block-diagonal form, where the characteristic polynomial of each block is a power of an irreducible polynomial. Since the Perron-Frobenius eigenvalue λ_{PF} has both geometric and algebraic multiplicity d, every eigenvalue algebraically conjugate to λ_{PF} also has multiplicity d. Thus there are d blocks whose characteristic polynomial has λ_{PF} for a root. We consider the constraints on the spectrum that can be obtained from these blocks alone.

Consider the projection of $\mathbf{k}\cdot\mathbf{L}_f/(2\pi)$ onto the large eigenspaces of the Perron-Frobenius block. Since only $d(b_{PF}+s_{PF})$ components of $\mathbf{t}$ (expressed in the new basis) contribute, this is the projection of $\mathbb{Q}^{d(b_{PF}+s_{PF})}$ onto $\mathbb{R}^{db_{PF}}$, whose real span is all of $\mathbb{R}^{db_{PF}}$. Intersected with the d-plane defined by a fixed $\mathbf{L}_f$, this gives a vector space of dimension at most $d(s_{PF}+1)$ in which $\mathbf{k}$ can live. □

5. One dimensional substitutions revisited

When discussing 1-dimensional substitutions, with n tile types $t_1, \ldots, t_n$, the conventional object of study is the *substitution matrix*, whose (i,j) entry gives the number of times that t_i appears in $\sigma(t_j)$. Indeed, in our previous study [CS] of one dimensional tilings, all the results were phrased in terms of eigenvalues and eigenspaces of the substitution matrix, rather than on the induced action of σ on homology. These results become much simpler when viewed homologically. In this section, the substitution matrix will be denoted M_s, while the matrix that gives the action of σ on a basis of recurrences will be denoted M_h.

In [CS] we defined, for each recurrence (z_1, z_2), a vector in $\mathbb{Z}^n$ that listed how many of each tile type appears in the (unique) path from z_1 to z_2. This vector $\mathbf{v}$ was called *full* if the vectors $(M_s)^k\mathbf{v}$, with k ranging from 0 to $n-1$, were linearly independent. Many of our theorems required the existence of a recurrence with a full vector. This is a strong condition, as it implies that $H_1(\Gamma)$ is a lattice of rank n. This is true when $n = 2$, or when the characteristic polynomial of M_s is irreducible, but is typically false for more complicated substitutions. When $H_1(\Gamma)$ has rank less than n, there are deformations of tile lengths that have no effect on the lengths of recurrences, and so lead to MLD tilings. By looking at M_h rather than M_s, we automatically avoid those extraneous modes.

Consider the difference between the following theorem, proved in [CS], and its restatement in terms of M_h:

Theorem 5.1 (CS)**.** *Suppose that all the eigenvalues of M_s are of magnitude 1 or greater, and that there exists a recurrence with a full vector. If the ratio of any two tile lengths is irrational, then the point spectrum is trivial.*

Theorem 5.2 (Corollary of Theorem 4.2)**.** *Suppose that all the eigenvalues of M_h are of magnitude 1 or greater. If the ratio of the lengths of any two recurrences is irrational, then the point spectrum is trivial.*

In addition to M_h not containing irrelevant information found in M_s, M_h may contain some relevant information *not* found in M_s. If $H_1(\Gamma^{(1)})$ has higher rank than $H_1(\Gamma^{(0)})$, then M_h contains information about the dynamical impact of changing the sizes of the collared tiles, and not merely the effect of changing the original, uncollared tiles.

As an example, consider the Thue-Morse substitution $(a \to ab, b \to ba)$, in which $M_s = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ has eigenvalues 2 and 0. Γ is the wedge of two circles, one representing the tile a and one representing the tile b, so $H_1(\Gamma) = \mathbb{Z}^2$, and the action of σ on $H_1(\Gamma)$ is described by M_s. However, $H_1(\Gamma^{(1)})$ has rank 3 [AP], and the eigenvalues of M_h are 2, -1, and 0. The additional large eigenvalue -1 shows that the dynamics of the Thue-Morse tiling space are in fact sensitive to changes in tile size. Changes in the size of the *uncollared* tiles have no qualitative effect, but changes in the size of the collared tiles (i.e., changes in tile size that depend on the local neighborhood of those tiles) can eliminate the point spectrum.

Acknowledgements. We thank Michael Baake, Charles Holton, Johannes Kellendonk, Brian Martensen, Charles Radin and Bob Williams for helpful discussions.

References

[AP] J. E. Anderson and I. F. Putnam, Topological invariants for substitution tilings and their associated C^*-algebras. *Ergodic Theory & Dynamical Systems* **18** (1998), 509-537.

[BBG] J. Bellissard, R. Benedetti, and J.-M. Gambaudo, Spaces of tilings, finite telescopic approximations and gap-labelling, to appear in *Commun. Math. Phys.*

[BG] R. Benedetti and G.-M. Gambaudo, On the dynamics of G-solenoids: applications to Delone sets, to appear in *Ergodic Theory & Dynamical Systems* .

[BSJ] M. Baake, M. Schlottmann and P. D. Jarvis, Quasiperiodic tilings with tenfold symmetry and equivalence with respect to local derivability. *J. Physics A* **24** (1991), 4637-4654.

[CS] A. Clark and L. Sadun, When size matters: subshifts and their related tiling spaces. *Ergodic Theory & Dynamical Systems*, to appear.

[G] F. Gähler, in preparation.

[K] J. Kellendonk, Pettern-equivariant functions and cohomology, *J. Phys. A* **36** (2003) 5765–5772.

[LM] D. Lind and B. Marcus, An Introduction to Symbolic Dynamics and Coding, Cambridge University Press, NY (1995).

[M1] B. Mossé, Puissances de mots et reconnaissabilité des points fixes d'une substitution. *Theoret. Comput. Sci.* **99** (1992), 327–334.

[M2] B. Mossé, Reconnaissabilité des substitutions et complexité des suites automatiques. *Bull. Soc. Math. France* **124** (1996), 329–346.

[ORS] N. Ormes, C. Radin, and L. Sadun, A homeomorphism invariant for substitution tiling spaces, *Geom. Dedicata* **90** (2002) 153–182.

[Pe] K. Petersen, Factor maps between tiling dynamical systems, *Forum Math.* **11** (1999) 503–512.

[Pr] N. M. Priebe, Towards a characterization of self-similar tilings in terms of derived Voronoi tesselations. *Geom. Dedicata* **79** (2000), 239–265.

[PS] N. Priebe and B. Solomyak, Characterization of planar pseudo-self-similar tilings. *Discrete Comput. Geom.* **26** (2001), 289–306.

[RS] C. Radin and L. Sadun, Isomorphism of hierarchical structures, *Ergodic Theory & Dynamical Systems* **21** (2001), 1239-1248.

[Sa] L. Sadun, Tiling spaces are inverse limits, preprint 2002, `math.DS/0210179`.

[So1] B. Solomyak, Dynamics of self-similar tilings. *Ergodic Theory & Dynamical Systems* **17** (1997), 695–738; Corrections in **19** (1999), 1685.

[So2] B. Solomyak, Non-periodicity implies unique composition for self-similar translationally-finite tilings. *Discrete Comput. Geom.* **20** (1998), 265–279.

[SW] L. Sadun and R. F. Williams, Tiling spaces are Cantor set fiber bundles. *Ergodic Theory & Dynamical Systems* **23** (2003), 307–316.

Alex Clark: Department of Mathematics, University of North Texas, Denton, Texas 76203

E-mail address: `AlexC@unt.edu`

Lorenzo Sadun: Department of Mathematics, The University of Texas at Austin, Austin, TX 78712-1082 U.S.A.

E-mail address: `sadun@math.utexas.edu`